Problem size of interest to scientists

RTFlame problem

The simulations Dean runs are at certain fixed problem sizes: Typically, 256**3 effective grid points and 512**3 effective grid points.

He will be running many 512**3 simulations on BG/P in the coming year. Currently he uses 4 racks of BG/P in VN mode (16,378 cores) for these simulations.

This number of cores is: a). Large enough so that the subsequent tree of 16**3 blocks fit in memory. b). Small enough so that there is sufficient work per process (~8-10 blocks / process).

He would like to be able to run the same 512**3 simulation on 32,768 cores. However, spreading the problem over this many cores results in ~4 blocks / process. Currently, the FLASH code does not scale well in this configuration. A possibility is poor load balance, but this has not been confirmed. So far these are computational issues that do not take into account IO. IO is a separate issue, and IO is found to perform very poorly at this scale. We do have a split IO mode which may alleviate the IO issues, but this has not been tested in RTFlame simulations.

WD_Def problem

The simulations Cal runs are at various problem sizes. In the following year he will be running simulations that use between 20,000-80,000 16**3 blocks. He is very interested in any potential optimisation of 50,000 block simulations.

The WD_Def weak scaling parameter files (attached) are used for simulations in which the ignition is at the star center. The subsequent evolution of this ignited region exercises the same physics as production simulations. (The difference in production runs is that the parameter files specify off-center ignition(s) which lead to rising flame bubble(s).)

In order for the same size simulations to finish in a reasonable time period on BG/P (i.e. equivalent to other supercomputers), we must spread a particular problem over more processors leading to ~6-8 blocks / process.

This would mean: 20,000 blocks on ~2,000-4,000 cores, 50,000 blocks on ~8,192 cores

In contrast, on other machines Cal can use ~30-50 blocks / process. This generally means that here the 20,000-80,000 block simulations are run in 1,000-4,000 processes.

RTFlame and WD_Def note

Through trial and error, we have found maxblocks=80 is a suitable maximum number of blocks for both RTFlame and WD_Def simulations when run in 512 MB per process.